Efficient Sum of Outer Products Dictionary Learning (SOUP-DIL) – The ℓ_0 Method

Saiprasad Ravishankar, *Member, IEEE,* Raj Rao Nadakuditi, *Member, IEEE,* and Jeffrey A. Fessler, *Fellow, IEEE*

arXiv:1511.08842v1 [cs.LG] 27 Nov 2015

***Abstract*—The sparsity of natural signals and images in a transform domain or dictionary has been extensively exploited in several applications such as compression, denoising and inverse problems. More recently, data-driven adaptation of synthesis dictionaries has shown promise in many applications compared to fixed or analytical dictionary models. However, dictionary learning problems are typically non-convex and NP-hard, and the usual alternating minimization approaches for these problems are often computationally expensive, with the computations dominated by the NP-hard synthesis sparse coding step. In this work, we investigate an efficient method for ℓ_0 "norm"-based dictionary learning by first approximating the training data set with a sum of sparse rank-one matrices and then using a block coordinate descent approach to estimate the unknowns. The proposed block coordinate descent algorithm involves efficient closed-form solutions. In particular, the sparse coding step involves a simple form of thresholding. We provide a convergence analysis for the proposed block coordinate descent approach. Our numerical experiments show the promising performance and significant speed-ups provided by our method over the classical K-SVD scheme in sparse signal representation and image denoising.**

***Index Terms*—Sparse representations, Dictionary learning, Image denoising, Fast algorithms, Machine learning, Convergence analysis, Block coordinate descent, Nonconvex optimization**

I. Introduction

The sparsity of natural signals and images in a transform domain or dictionary has been extensively exploited in several applications such as compression [1], denoising, compressed sensing and other inverse problems. Well-known models for sparsity include the synthesis, analysis [2], and the transform [3], [4] (or generalized analysis) models. More recently, the data-driven adaptation of sparse signal models has benefited many applications [4]–[12] compared to fixed or analytical models. In this work, we focus on the data-driven adaptation of the synthesis model and present a simple and highly efficient algorithm with convergence analysis and applications. In the following, we first briefly review the topic of synthesis dictionary learning before summarizing the contributions of this work.

A. Synthesis Model and Dictionary Learning

The well-known synthesis model suggests that a signal $y \in \mathbb{R}^n$ is approximately a linear combination of a small subset of atoms or columns of a dictionary $D \in \mathbb{R}^{n \times J}$, i.e., $y = Dx + e$ with $x \in \mathbb{R}^J$ sparse, and e is assumed to be a small modeling error or approximation error in the signal domain [13]. We say that $x \in \mathbb{R}^J$ is sparse if $\|x\|_0 \ll n$, where the ℓ_0 "norm" counts the number of non-zero entries in x. Since different candidate signals may be approximately spanned by different subsets of columns in the dictionary D, the synthesis model is also known as a union of subspaces model [14], [15]. When $n = J$ and D is full rank, it is a basis. Else when $J > n$, D is called an overcomplete dictionary. Because of their richness, overcomplete dictionaries can provide highly sparse (i.e., with few non-zeros) representations of data and are popular.

For a given signal y and dictionary D, the process of finding a sparse representation x involves solving the well-known synthesis sparse coding problem. This problem is to minimize $\|y - Dx\|_2^2$ subject to $\|x\|_0 \leq s$, where s is some set sparsity level. The synthesis sparse coding problem is NP-hard (Non-deterministic Polynomial-time hard) [16], [17]. Numerous algorithms [18]–[25] including greedy and relaxation algorithms have been proposed for this problem. While some of these algorithms are guaranteed to provide the correct solution under certain conditions, these conditions are often restrictive and violated in applications. Moreover, these sparse coding algorithms typically tend to be computationally expensive for large-scale problems.

More recently, the data-driven adaptation of synthesis dictionaries, called dictionary learning, has been investigated in several works [8], [26]–[29]. Dictionary learning has shown promise in several applications including compression, denoising, inpainting, deblurring, demosaicing, super-resolution, and classification [5], [6], [30]–[37]. It has also been demonstrated to be useful in inverse problems such as those in tomography [38], and magnetic resonance imaging (MRI) [9], [39]–[41].

Given a collection of training signals $\{y_i\}_{i=1}^N$ that are represented as columns of the matrix $Y \in \mathbb{R}^{n \times N}$, the dictionary learning problem is often formulated as follows [28]

$$\text{(P0)} \ \min_{D,X} \|Y - DX\|_F^2 \ \text{ s.t. } \ \|x_i\|_0 \leq s \ \forall i, \ \|d_j\|_2 = 1 \, \forall j.$$

Here, d_j and x_i denote the columns of the dictionary $D \in \mathbb{R}^{n \times J}$ and sparse code matrix $X \in \mathbb{R}^{J \times N}$, respectively, and s denotes the maximum sparsity level (non-zeros in representations x_i) allowed for each training signal. The columns of the dictionary are constrained to have unit norm to avoid the scaling ambiguity [42]. Variants of Problem (P0) include replacing the ℓ_0 "norm" for sparsity with an ℓ_1 norm or an alternative sparsity criterion, or enforcing additional properties

This work was supported in part by the following grants: ONR grant N00014-15-1-2141, DARPA Young Faculty Award D14AP00086, ARO MURI grants W911NF-11-1-0391 and 2015-05174-05, NIH grant U01 EB01875301, and a UM-SJTU seed grant.

S. Ravishankar, R. Nadakuditi, and J. Fessler are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, 48109 USA emails: (ravisha, rajnrao, fessler)@umich.edu.

(e.g., incoherence [35], [43]) for the dictionary D, or solving an online version (where the dictionary is updated sequentially as new training signals arrive) of the problem [8].

Algorithms for Problem (P0) or its variants [8], [27]–[29], [44]–[52] typically alternate in some form between a *sparse coding step* (updating X), and a *dictionary update step* (solving for D). Some of these algorithms (e.g., [28], [48], [50]) also partially update X in the dictionary update step. A few recent methods attempt to solve for D and X jointly in an iterative fashion [53], [54]. The K-SVD method [28] has been particularly popular in numerous applications [5]–[7], [9], [30], [39]. However, Problem (P0) is highly non-convex and NP-hard, and most dictionary learning approaches lack proven convergence guarantees. Moreover, the algorithms for (P0) tend to be computationally expensive (particularly alternating-type algorithms), with the computations usually dominated by the synthesis sparse coding step.

Some recent works [52], [55]–[59] have studied the convergence of (specific) synthesis dictionary learning algorithms. However, these dictionary learning methods have not been demonstrated to be useful in applications such as image denoising. Bao et al. [52] find that their method, although a fast proximal scheme, denoises less effectively than the K-SVD method [5]. Many prior works use restrictive assumptions (e.g., noiseless data, etc.) for their convergence results.

B. Contributions

This work focuses on synthesis dictionary learning and investigates a dictionary learning problem with an ℓ_0 sparsity penalty instead of constraints. Our approach first models the training data set as an approximate sum of sparse rank-one matrices (or outer products). Then, we use a simple and exact block coordinate descent approach to estimate the factors of the various rank-one matrices. In particular, the sparse coding step in the algorithm uses a simple form of thresholding, and the dictionary atom update step involves a matrix-vector product that is computed efficiently. We provide a convergence analysis of the proposed block coordinate descent method. Our numerical experiments demonstrate the promising performance and significant speed-ups provided by our method over the classicial K-SVD dictionary learning scheme in sparse signal representation and image denoising. Importantly, our algorithm is similar to just the dictionary update step of K-SVD, and yet performs comparably or much better and faster than K-SVD in applications.

Our work shares similarities with a recent dictionary learning approach [60] that exploits a sum of outer products model for the training data. However, the specific problem formulation and algorithm studied in this work differ from the prior work [60]. Importantly, unlike the previous approach [60], we provide a detailed convergence analysis of the proposed algorithm and also demonstrate its usefulness (both in terms of speed and quality of results) in sparse representation of natural signals and in image denoising. The sum of outer products model and related algorithm can accomodate various structures or constraints on the dictionary or sparse codes.

C. Organization

The rest of this paper is organized as follows. Section II discusses the problem formulation for dictionary learning, along with some potential alternatives. Section III presents our dictionary learning algorithm and discusses its computational properties. Section IV presents a convergence analysis for the proposed algorithm. Section V illustrates the empirical convergence behavior of our method and demonstrates its usefulness for sparse signal representation and image denoising. Section VI presents conclusions and proposals for future work.

II. Problem Formulations

A. ℓ_0 Penalized Formulation

We consider a sparsity penalized variant of Problem (P0) [52]. Specifically, by replacing the sparsity constraints in (P0) with an ℓ_0 penalty $\sum_{i=1}^{N} \|x_i\|_0$ and introducing a variable $C = X^T \in \mathbb{R}^{N \times J}$, where $(\cdot)^T$ denotes the matrix transpose operation, we arrive at the following formulation:

$$\min_{D,C} \left\|Y - DC^T\right\|_F^2 + \lambda^2 \|C\|_0 \ \text{ s.t. } \ \|d_j\|_2 = 1 \, \forall \, j. \quad (1)$$

where $\|C\|_0$ counts the number of non-zeros in the matrix C, and $\lambda^2 > 0$ is a weight for the sparsity penalty that controls the number of non-zeros in the sparse representation.

Next, we express the matrix DC^T in (1) as a sum of (sparse) rank-one matrices or outer products $\sum_{j=1}^{J} d_j c_j^T$, where c_j is the jth column of C. This is a natural representation of the training data Y because it separates out the contributions of the various atoms in representing the data. It may also provide a natural way to set the number of atoms (degrees of freedom) in the dictionary. In particular, atoms of a dictionary whose contributions to the data (Y) representation error or modeling error are small could be dropped. Such a Sum of OUter Products (SOUP) representation has been exploited in previous dictionary learning algorithms [28], [48]. With this model, we rewrite (1) as follows, where $\|C\|_0 = \sum_{j=1}^{J} \|c_j\|_0$:

$$\text{(P1)} \ \min_{\{d_j, c_j\}} \left\|Y - \sum_{j=1}^{J} d_j c_j^T\right\|_F^2 + \lambda^2 \sum_{j=1}^{J} \|c_j\|_0$$
$$\text{s.t. } \ \|d_j\|_2 = 1, \, \|c_j\|_\infty \leq L \, \forall \, j.$$

As in Problem (P0), the matrix $d_j c_j^T$ in (P1) is invariant to joint scaling of d_j and c_j as αd_j and $(1/\alpha) c_j$, for $\alpha \neq 0$. The constraint $\|d_j\|_2 = 1$ helps remove this scaling ambiguity. We also enforce the constraint $\|c_j\|_\infty \leq L$, with $L > 0$, in (P1) [52] (e.g., $L = \|Y\|_F$). Consider a dictionary D that has a column d_j that repeats. Then, in this case, the outer product expansion of Y in (P1) could have both the terms $d_j c_j^T$ and $-d_j c_j^T$ with c_j that is highly sparse, and the objective would be invariant to (arbitrarily) large scalings[1] of c_j (i.e., non-coercive objective). The ℓ_∞ constraints on the columns of C (that constrain the magnitudes of entries of C) prevent such a scenario. Problem (P1) is designed to learn the factors $\{d_j\}_{j=1}^{J}$ and $\{c_j\}_{j=1}^{J}$ that enable the best SOUP sparse representation of Y.

[1] Such degenerate representations of Y, however, cannot be minimizers in the problem because they simply increase the ℓ_0 sparsity penalty without affecting the fitting error (the first term) in the cost.

B. Alternative Formulations

Several variants of the SOUP learning Problem (P1) could be constructed, with interesting effects. For example, the ℓ_0 "norm" for sparsity could be replaced by the ℓ_1 norm [60]. Another interesting alternative to (P1) involves enforcing p-block-orthogonality constraints on the dictionary D. The dictionary in this case is split into blocks (instead of individual atoms), each of which has p atoms that are orthogonal to each other. For $p = 2$, the constraints take the form $d_{2j-1}^T d_{2j} = 0$, $1 \leq j \leq J/2$. In the extreme (more constrained) case of $p = n$, the dictionary would be made of several (square) orthonormal blocks. For tensor-type data, Problem (P1) can be modified by enforcing the dictionary atoms to be in the form of a Kronecker product. The algorithm proposed in Section III can be easily extended to accomodate such variants of Problem (P1). We do not fully explore such alternatives in this work due to space constraints, and a more detailed investigation of these will be presented elsewhere.

III. Algorithm and Properties

A. Algorithm

We propose a block coordinate descent method to estimate the unknown variables in Problem (P1). For each j ($1 \leq j \leq J$), our algorithm has two key steps. First, we solve (P1) with respect to c_j keeping all the other variables fixed. We refer to this step as the *sparse coding step* in our method. Once c_j is updated, we solve (P1) with respect to d_j keeping all other variables fixed. This step is referred to as the *dictionary atom update step* or simply *dictionary update step*. The algorithm updates the factors of the various rank-one matrices one-by-one. We next describe the sparse coding and dictionary atom update steps.

1) Sparse Coding Step: Minimizing (P1) with respect to c_j leads to the following non-convex problem, where $E_j \triangleq Y - \sum_{k \neq j} d_k c_k^T$ is a fixed matrix based on the most recent values of all other atoms and coefficients:

$$\min_{c_j} \ \left\| E_j - d_j c_j^T \right\|_F^2 + \lambda^2 \left\| c_j \right\|_0 \ \text{ s.t. } \ \left\| c_j \right\|_\infty \leq L. \tag{2}$$

The following proposition provides the solution to Problem (2), where the hard-thresholding operator $H_\lambda(\cdot)$ is defined as

$$(H_\lambda(b))_i = \begin{cases} 0, & |b_i| < \lambda \\ b_i, & |b_i| \geq \lambda \end{cases} \tag{3}$$

with $b \in \mathbb{R}^N$, and the subscript i indexes vector entries. We assume that the bound $L > \lambda$ and let 1_N denote a vector of ones of length N. The operation "$\odot$" denotes element-wise multiplication, $\text{sign}(\cdot)$ computes the signs of the elements of a vector, and $z = \min(a, u)$ for vectors $a, u \in \mathbb{R}^N$ denotes the element-wise minimum operation, i.e., $z_i = \min(a_i, b_i)$, $1 \leq i \leq N$.

Proposition 1: Given $E_j \in \mathbb{R}^{n \times N}$ and $d_j \in \mathbb{R}^n$, and assuming $L > \lambda$, a global minimizer of the sparse coding problem (2) is obtained by the following truncated hard-thresholding operation:

$$\hat{c}_j = \min\left(\left| H_\lambda\left(E_j^T d_j \right) \right|, L 1_N \right) \odot \text{sign}\left(H_\lambda\left(E_j^T d_j \right) \right) \tag{4}$$

The minimizer of (2) is unique if and only if the vector $E_j^T d_j$ has no entry with a magnitude of λ.

Proof: First, for a vector d_j that has unit ℓ_2 norm, we have the following equality

$$\begin{aligned} \left\| E_j - d_j c_j^T \right\|_F^2 &= \left\| E_j \right\|_F^2 + \left\| c_j \right\|_2^2 - 2 c_j^T E_j^T d_j \\ &= \left\| c_j - E_j^T d_j \right\|_2^2 + \left\| E_j \right\|_F^2 - \left\| E_j^T d_j \right\|_F^2 \end{aligned} \tag{5}$$

By substituting (5) into (2), it is clear that Problem (2) is equivalent to

$$\min_{c_j} \ \left\| c_j - E_j^T d_j \right\|_2^2 + \lambda^2 \left\| c_j \right\|_0 \ \text{ s.t. } \ \left\| c_j \right\|_\infty \leq L. \tag{6}$$

Define $b \triangleq E_j^T d_j$. Then, the objective in (6) simplifies to $\sum_{i=1}^N \left\{ |c_{ji} - b_i|^2 + \lambda^2 \, \theta(c_{ji}) \right\}$ with

$$\theta(a) = \begin{cases} 0, & \text{if } \ a = 0 \\ 1, & \text{if } \ a \neq 0 \end{cases} \tag{7}$$

Therefore, we solve for each entry c_{ji} of c_j as

$$\hat{c}_{ji} = \arg\min_{c_{ji}} |c_{ji} - b_i|^2 + \lambda^2 \, \theta(c_{ji}) \ \text{ s.t. } \ |c_{ji}| \leq L. \tag{8}$$

It is straightforward to show that when $|b_i| \leq L$ (case (a)),

$$\hat{c}_{ji} = \begin{cases} 0, & \text{if } \ b_i^2 < \lambda^2 \\ b_i, & \text{if } \ b_i^2 > \lambda^2 \end{cases} \tag{9}$$

When $|b_i| = \lambda$ ($\lambda < L$), the optimal $\hat{c}_{ji}$ can be either b_i or 0 (non-unique), and both these settings achieve the minimum objective value λ^2 in (8). Next, when $|b_i| > L$ (case (b)), we have

$$\hat{c}_{ji} = \begin{cases} 0, & \text{if } \ b_i^2 < \left(L \, \text{sign}(b_i) - b_i \right)^2 + \lambda^2 \\ L \, \text{sign}(b_i), & \text{if } \ b_i^2 > \left(L \, \text{sign}(b_i) - b_i \right)^2 + \lambda^2 \end{cases} \tag{10}$$

Since $L > \lambda$, clearly $b_i^2 > \left(L \, \text{sign}(b_i) - b_i \right)^2 + \lambda^2$ in (10).

Thus, an optimal $\hat{c}_{ji}$ in (8) is compactly written as $\hat{c}_{ji} = \min\left(\left| H_\lambda\left(b_i \right) \right|, L \right) \cdot \text{sign}\left(b_i \right)$, thereby establishing (4). The condition for uniqueness of the sparse coding solution follows from the arguments for case (a) above. ■

2) Dictionary Atom Update Step: Minimizing (P1) with respect to d_j leads to the following non-convex problem:

$$\min_{d_j} \ \left\| E_j - d_j c_j^T \right\|_F^2 \ \text{ s.t. } \ \left\| d_j \right\|_2 = 1 \tag{11}$$

Proposition 2 provides the closed-form solution for (11).

Proposition 2: Given $E_j \in \mathbb{R}^{n \times N}$ and $c_j \in \mathbb{R}^N$, a global minimizer of the dictionary atom update problem (11) is

$$\hat{d}_j = \begin{cases} \frac{E_j c_j}{\left\| E_j c_j \right\|_2}, & \text{if } \ c_j \neq 0 \\ v_1, & \text{if } \ c_j = 0 \end{cases} \tag{12}$$

where v_1 is the first column of the $n \times n$ identity matrix. The solution is unique if and only if $c_j \neq 0$.

Proof: First, for a vector d_j that has unit ℓ_2 norm, the following holds:

$$\left\| E_j - d_j c_j^T \right\|_F^2 = \left\| E_j \right\|_F^2 + \left\| c_j \right\|_2^2 - 2 \, d_j^T E_j c_j \tag{13}$$

SOUP-DIL Algorithm

Inputs : Training data $Y \in \mathbb{R}^{n \times N}$, weight λ, upper bound L, and number of iterations K.

Outputs : Columns $\left\{d_j^K\right\}_{j=1}^{J}$ of the learned dictionary, and the learned sparse coefficients $\left\{c_j^K\right\}_{j=1}^{J}$.

Initial Estimates: $\left\{d_j^0, c_j^0\right\}_{j=1}^{J}$. (Often $c_j^0 = 0 \ \forall \ j$.)

For t = $1 : K$ **Repeat**

For j = $1 : J$ **Repeat**

1) $C = \left[c_1^t, ..., c_{j-1}^t, c_j^{t-1}, ..., c_J^{t-1}\right]$.
$D = \left[d_1^t, ..., d_{j-1}^t, d_j^{t-1}, ..., d_J^{t-1}\right]$.

2) **Sparse coding:**

$$b^t = Y^T d_j^{t-1} - CD^T d_j^{t-1} + c_j^{t-1} \tag{17}$$

$$c_j^t = \min\left(\left|H_\lambda\left(b^t\right)\right|, L1_N\right) \odot \operatorname{sign}\left(H_\lambda\left(b^t\right)\right) \tag{18}$$

3) **Dictionary atom update:**

$$h^t = Yc_j^t - DC^T c_j^t + d_j^{t-1}\left(c_j^{t-1}\right)^T c_j^t \tag{19}$$

$$d_j^t = \begin{cases} \frac{h^t}{\|h^t\|_2}, & \text{if } c_j^t \neq 0 \\ v_1, & \text{if } c_j^t = 0 \end{cases} \tag{20}$$

End

End

Fig. 1. The SOUP-DIL Algorithm (or SOUP-DILLO Algorithm, due to the use of the ℓ_0 "norm") for Problem (P1). Superscript of t denotes the iterates in the algorithm. Although we perform the sparse coding step prior to the dictionary update step, one could also potentially switch this order. The vectors b^t and h^t above can be computed very efficiently via sparse operations.

Upon substituting (13) into (11), Problem (11) simplifies to

$$\max_{d_j} \ d_j^T E_j c_j \quad \text{s.t.} \quad \|d_j\|_2 = 1. \tag{14}$$

By the Cauchy Schwarz inequality $d_j^T E_j c_j \leq \|E_j c_j\|_2$ for unit norm d_j. Thus, a solution to (14) that achieves the value $\|E_j c_j\|_2$ for the objective is

$$\hat{d}_j = \begin{cases} \frac{E_j c_j}{\|E_j c_j\|_2}, & \text{if } E_j c_j \neq 0 \\ v_1, & \text{if } E_j c_j = 0 \end{cases} \tag{15}$$

Obviously, any $d \in \mathbb{R}^n$ would be a minimizer (non-unique) in (14) when $E_j c_j = 0$. In particular $\hat{d}_j = v_1$ works.

Next, we show that $E_j c_j = 0$ in our algorithm if and only if $c_j = 0$. This result together with (15) immediately establishes the proposition. Since the c_j used in the dictionary atom update step (11) was obtained as a minimizer in the preceding sparse coding step (2), we have the following inequality for all $c \in \mathbb{R}^N$ with $\|c\|_\infty \leq L$, and $\tilde{d}_j$ denotes the jth atom in the preceding sparse coding step:

$$\left\|E_j - \tilde{d}_j c_j^T\right\|_F^2 + \lambda^2 \|c_j\|_0 \leq \left\|E_j - \tilde{d}_j c^T\right\|_F^2 + \lambda^2 \|c\|_0 \tag{16}$$

If $E_j c_j = 0$, the left hand side above simplifies to $\|E_j\|_F^2 + \|c_j\|_2^2 + \lambda^2 \|c_j\|_0$, which is clearly minimal when $c_j = 0$. Thus, when $E_j c_j = 0$, we must also have $c_j = 0$. ■

While Propositions 1 and 2 provide the minimizers of (2) and (11) for the case of real-valued matrices/vectors in the problems, these solutions are trivially extended to the complex-valued case (that may be useful in applications such as magnetic resonance imaging [9]) by using a Hermitian transpose.

Fig. 1 shows the Sum of OUter Products DIctionary Learning (SOUP-DIL) Algorithm (or SOUP-DILLO Algorithm, due to the use of the ℓ_0 "norm") for Problem (P1). The algorithm assumes that an initial estimate $\left\{d_j^0, c_j^0\right\}_{j=1}^{J}$ for the variables is provided. For example, the initial sparse coefficients could be set to zero, and the initial dictionary could be a known analytical dictionary such as the overcomplete DCT [5]. When $c_j^t = 0$, setting $d_j^t = v_1$ in (20) in the algorithm could also be replaced with other (equivalent) settings such as $d_j^t = d_j^{t-1}$ or setting d_j^t to a random unit norm vector. All of these settings have been observed to work well in practice. A random ordering of the atom/sparse coefficient updates in Fig. 1 (i.e., random j sequence) also helps in practice (in accelerating convergence) compared to cycling in the same order 1 through J every iteration. One could also alternate several times between the sparse coding and dictionary atom update steps for each j in Fig. 1. However, this would increase computation.

B. Computational Cost

For each iteration t in Fig. 1, our algorithm involves J sparse code and dictionary atom updates. The sparse coding and atom update steps involve matrix-vector products for computing b^t and h^t, respectively. An alternative approach to the one in Fig. 1 involves computing the matrix $E_j^t = Y - \sum_{k<j} d_k^t \left(c_k^t\right)^T - \sum_{k>j} d_k^{t-1}\left(c_k^{t-1}\right)^T$ (as in Propositions 1 and 2) directly at the beginning of each inner j iteration. This matrix can be updated sequentially and efficiently for each j by adding and subtracting appropriate sparse rank-one matrices. However, this alternative approach requires storing $E_j^t \in \mathbb{R}^{n \times N}$, which is often a large matrix for large N and n. Instead, the procedure adopted by us in Fig. 1 helps save memory usage. We now discuss the cost of each sparse coding and atom update procedure in Fig. 1.

Consider the tth iteration and the jth inner iteration in Fig. 1, consisting of the update of the jth dictionary atom and its corresponding sparse coefficients. As in Fig. 1, let $D \in \mathbb{R}^{n \times J}$ be the dictionary whose columns are the current estimates of the atoms (at the start of the jth inner iteration), and let $C \in \mathbb{R}^{N \times J}$ be the corresponding sparse coefficients matrix. (The index t on D and C is dropped to keep the notation simple.) Assume that the matrix C has $\alpha N n$ non-zeros, with $\alpha \ll 1$ typically. This translates to an average of αn non-zeros per row of C or $\alpha N n / J$ non-zeros per column of C. We refer to α as the sparsity factor of C.

The sparse coding step involves computing the right hand side of (17). While computing $Y^T d_j^{t-1}$ in (17) requires Nn multiply-add operations, computing $CD^T d_j^{t-1}$ using matrix-vector products requires $Jn + \alpha N n$ multiply-add operations. Computing the difference of these vectors and summing that result with the sparse c_j^{t-1} requires less than $2N$ additions. The hard thresholding-type operation (18) to obtain c_j^t requires at most $2N$ comparisons.

Next, when $c_j^t \neq 0$, the dictionary atom update step requires computing the right hand side of (19). Since c_j^t is sparse with say r_j non-zeros, computing Yc_j^t in (19) requires nr_j multiply-add operations, and computing $DC^T c_j^t$ using matrix-vector products requires less than $Jn + \alpha Nn$ multiply-add operations. The cost of the remaining operations in (19) and for normalizing h^t (20) is negligible.

Thus, the net cost of the $J \geq n$ inner iterations in iteration t in Fig. 1 is dominated (for $N \gg J, n$) by $NJn + 2\alpha_m NJn + \beta Nn^2$, where α_m is the maximum sparsity factor of the estimated C's during the inner iterations, and β is the sparsity factor of the estimated C at the end of iteration t. Thus, the cost per iteration of the block coordinate descent SOUP-DIL Algorithm is about $(1 + \alpha')NJn$, with $\alpha' \ll 1$ typically. On the other hand, the proximal alternating algorithm proposed very recently[2] by Bao et al. [52], [61] has a per-iteration computational cost of at least $2NJn + 6\alpha NJn + 4\alpha Nn^2$. This is clearly more computation than SOUP-DIL.

Assuming $J \propto n$, the cost per iteration of the SOUP-DIL Algorithm scales as $O(Nn^2)$. This is lower than the per-iteration cost of learning an $n \times J$ synthesis dictionary D using K-SVD [28], which scales[3] (assuming that the synthesis sparsity level $s \propto n$ and $J \propto n$ in K-SVD) as $O(Nn^3)$.

As illustrated in Section V-B, our algorithm converges in few iterations in practice. Therefore, the per-iteration computational advantages also translate to net computational advantages in practice. The low computational cost of our approach could be particularly useful for big data applications, or higher dimensional (3D or 4D) applications.

IV. Convergence Analysis

This section presents a convergence analysis for the proposed SOUP-DIL Algorithm for Problem (P1). Problem (P1) is highly non-convex due to the ℓ_0 penalty for sparsity, the unit ℓ_2 norm constraints on atoms of D, and the term $\left\| Y - \sum_{j=1}^{J} d_j c_j^T \right\|_F^2$ that is a non-convex function involving the products of multiple unknown vectors. The proposed algorithm is an exact block coordinate descent procedure for Problem (P1). However, due to the high degree of non-convexity, standard results on convergence of block coordinate descent methods (e.g., [63]) do not apply here. More recent works [64] on the convergence of block coordinate descent schemes use assumptions (such as multi-convexity) that do not hold in our setting. Here, we discuss the convergence of our algorithm to the critical points in the problem. In the following, we first present some definitions and notations, before stating the main convergence results.

A. Definitions and Notations

First, we review the Fréchet sub-differential of a function [65], [66]. The norm and inner product notation in Definition 2 correspond to the euclidean ℓ_2 settings.

[2]This method also involves more parameters than our scheme.

[3]When $s \propto n$ and $J \propto n$, the per-iteration computational cost of the efficient implementation of K-SVD [62] also scales similarly as $O(Nn^3)$.

Definition 1: For a function $g : \mathbb{R}^p \mapsto (-\infty, +\infty]$, its domain is defined as $\mathrm{dom} g = \{x \in \mathbb{R}^p : g(x) < +\infty\}$. Function g is proper if $\mathrm{dom} g$ is nonempty.

Definition 2: Let $g : \mathbb{R}^p \mapsto (-\infty, +\infty]$ be a proper function and let $x \in \mathrm{dom} g$. The Fréchet sub-differential of the function g at x is the following set denoted as $\hat{\partial} g(x)$:

$$\left\{ h \in \mathbb{R}^p : \liminf_{b \to x, b \neq x} \frac{1}{\|b - x\|} \left(g(b) - g(x) - \langle b - x, h \rangle \right) \geq 0 \right\}$$

If $x \notin \mathrm{dom} g$, then $\hat{\partial} g(x) \triangleq \emptyset$, the empty set. The sub-differential of g at x is the set $\partial g(x)$ defined as

$$\left\{ \tilde{h} \in \mathbb{R}^p : \exists x_k \to x, g(x_k) \to g(x), h_k \in \hat{\partial} g(x_k) \to \tilde{h} \right\}.$$

A necessary condition for $x \in \mathbb{R}^p$ to be a minimizer of the function g is that x is a *critical point* of g, i.e., $0 \in \partial g(x)$. Critical points are considered to be "generalized stationary points" [65].

We say that a sequence $\{z^t\} \subset \mathbb{R}^p$ has an accumulation point z, if there is a subsequence that converges to z.

The constraints $\|d_j\|_2 = 1$, $1 \leq j \leq J$, in (P1) can instead be added as penalties in the cost by using barrier functions $\chi(d_j)$ (taking the value $+\infty$ when the norm constraint is violated, and is zero otherwise). The constraints $\|c_j\|_\infty \leq L$, $1 \leq j \leq J$, can also be similarly replaced with barrier penalties $\psi(c_j)$. Then, we rewrite (P1) in unconstrained form with the following objective:

$$f(C, D) = f(c_1, c_2, ..., c_J, d_1, d_2, ..., d_J) = \lambda^2 \sum_{j=1}^{J} \|c_j\|_0 + \left\| Y - \sum_{j=1}^{J} d_j c_j^T \right\|_F^2 + \sum_{j=1}^{J} \chi(d_j) + \sum_{j=1}^{J} \psi(c_j) \qquad (21)$$

For the SOUP-DIL Algorithm, the iterates computed in the tth outer iteration are denoted by the 2J-tuple $(c_1^t, d_1^t, c_2^t, d_2^t, ..., c_J^t, d_J^t)$, or alternatively by the pair of matrices (C^t, D^t).

B. Main Results

Assume that the initial (C^0, D^0) satisfies the constraints in (P1). We then have the following simple monotonicity and consistency result.

Theorem 1: Let $\{C^t, D^t\}$ denote the iterate sequence generated by the SOUP-DIL Algorithm with training data $Y \in \mathbb{R}^{n \times N}$ and initial (C^0, D^0). Then, the objective sequence $\{f^t\}$ with $f^t \triangleq f(C^t, D^t)$ is monotone decreasing, and converges to a finite value, say f^*. Moreover, the iterate sequence $\{C^t, D^t\}$ is bounded, and all its accumulation points are equivalent in the sense that they achieve the exact same value f^* of the objective.

Proof: See Appendix A. ■

Theorem 1 establishes that for each initial point (C^0, D^0), the bounded iterate sequence in the SOUP-DIL Algorithm is such that all its accumulation points achieve the same value f^* of the objective. They are equivalent in that sense. The value of f^* could vary with different initalizations. We thus have the following corollary of Theorem 1 that holds because

the distance between a bounded sequence and its (non-empty and compact) set of accumulation points converges to zero.

Corollary 1: For each (C^0, D^0), the iterate sequence in the SOUP-DIL Algorithm converges to an equivalence class of accumulation points.

The following Theorem 2 considers a very special case of SOUP dictionary learning, where the dictionary has a single atom. In this case, the SOUP learning Problem (P1) is the problem of obtaining a sparse rank-one approximation of the training matrix Y. In this case, Theorem 2 establishes that the iterates in the algorithm converge to the set of critical points (i.e., the distance between the iterates and the set converges to zero) of the objective f.

Theorem 2: Consider the SOUP-DIL Algorithm with $J = 1$. Let $\{c^t, d^t\}$ denote the bounded iterate sequence generated by the algorithm in this case with data $Y \in \mathbb{R}^{n \times N}$ and initial (c^0, d^0). Then, every accumulation point of the iterate sequence is a critical point of the objective f, i.e., the iterates converge to the set of critical points of f.

Proof: See Appendix B. ■

For the general case ($J \neq 1$), we have the following result.

Theorem 3: Let $\{C^t, D^t\}$ denote the bounded iterate sequence generated by the SOUP-DIL Algorithm with training data $Y \in \mathbb{R}^{n \times N}$ and initial (C^0, D^0). Suppose each accumulation point (C, D) of the iterate sequence is such that the matrix B with columns $b_j = E_j^T d_j$ and $E_j = Y - DC^T + d_j c_j^T$, has no entry with magnitude λ. Then every accumulation point of the iterate sequence is a critical point of the objective $f(C, D)$. Moreover, the two sequences with terms $\left\|D^t - D^{t-1}\right\|_F$ and $\left\|C^t - C^{t-1}\right\|_F$ respectively, both converge to zero.

Proof: See Appendix C. ■

Theorem 3 establishes that the iterates in SOUP-DIL converge to the set of critical points of $f(C, D)$. For each initial (C^0, D^0), the iterate sequence in the algorithm converges (using Corollary 1) to an equivalence class of critical points of f. Theorem 3 also establishes that $\left\|D^t - D^{t-1}\right\|_F \to 0$ and $\left\|C^t - C^{t-1}\right\|_F \to 0$, thereby implying that the sparse approximation to the training data $Z^t = D^t \left(C^t\right)^T$ is such that $\left\|Z^t - Z^{t-1}\right\|_F \to 0$. These are necessary but not sufficient conditions for the convergence of the entire sequences $\{D^t\}$, $\{C^t\}$, and $\{Z^t\}$. The assumption on the entries of the matrix B in Theorem 3 is equivalent to assuming that for every $1 \leq j \leq J$, there is a unique minimizer of f with respect to c_j with all other variables fixed to their values in the accumulation point (C, D).

Although Theorem 3 uses a uniqueness condition, the following conjecture postulates that provided the following Assumption 1 (that uses a probabilistic model for the data) holds, the uniqueness condition holds with probability 1, i.e., the probability of a tie in assigning sparse codes is zero.

Assumption 1. The training signals $y_i \in \mathbb{R}^n$ for $1 \leq i \leq N$, are drawn independently from an absolutely continuous probability measure over the n-dimensional ball $S \triangleq \{y \in \mathbb{R}^n : \|y\|_2 \leq \beta_0\}$ for some $\beta_0 > 0$.

Conjecture 1: Let Assumption 1 hold. Then, with probability 1, every accumulation point (C, D) of the iterate sequence in the SOUP-DIL Algorithm is such that for each $1 \leq j \leq J$, the minimizer of $f(c_1, ..., c_{j-1}, \tilde{c}_j, c_{j+1}, ..., c_J, d_1, ..., d_J)$ with respect to $\tilde{c}_j$ is unique.

If Conjecture 1 holds, then every accumulation point of the iterate sequence in the SOUP-DIL Algorithm is immediately a critical point of $f(C, D)$ with probability 1.

V. Numerical Experiments

A. Framework

This section presents numerical results illustrating the practical convergence behavior of the proposed algorithm, as well as its usefulness in sparse signal representation and image denoising. For demonstrating the convergence behavior and quality of sparse signal representations, we consider data formed using vectorized 2D patches of natural images. The SOUP-DIL Algorithm is used to learn sparse representations for such data, and we study the signal representation ability of the proposed algorithm at various values of λ (i.e., various sparsities). We also present results for patch-based image denoising using Problem (P1).

Our dictionary learning implementation was coded in Matlab version R2015a. We compare the performance of dictionaries learned using our method to those learned using the classical K-SVD[4] method [5], [28]. For K-SVD learning and denoising [5], [28], we consider the original Matlab implementation of the methods available from Michael Elad's website [68]. A fast version[5] of K-SVD [62] that also uses MEX/C implementations of sparse coding and some sub-steps of dictionary update, is publicly available [69]. The Matlab implementation of our method is not currently optimized for efficiency. Therefore, we compare the runtimes achieved by our unoptimized Matlab implementation to both the original Matlab and the efficient (partial MEX/C) implementations of K-SVD for a fair and instructive comparison. For K-SVD, we used the built-in parameter settings of the author's implementations, unless otherwise stated. All computations were performed with an Intel Xeon CPU X3230 at 2.66 GHz and 8 GB memory, employing a 64-bit Windows 7 operating system.

For sparse representation of data Y, we use the *normalized sparse representation error* (NSRE) $\left\|Y - DC^T\right\|_F / \|Y\|_F$ to measure the performance of the learned dictonaries. For image denoising, similar to prior work, we measure the peak-signal-to-noise ratio (PSNR) computed between the true noiseless reference and the noisy or denoised images.

B. Convergence Experiment

To study the convergence behavior of the proposed SOUP-DIL Algorithm, we extracted 3×10^4 patches of size 8×8 from randomly chosen locations in the images Barbara, Boat,

[4]The K-SVD method is a highly popular dictionary learning scheme that has been applied to many image processing applications including denoising [5], [6] and MR image reconstruction [9]. Mairal et al. [33] proposed a non-local method for image denoising that also exploits learned dictionaries and achieves denoising performance comparable to the well-known BM3D [67] denoising method. Similar extensions to our proposed method for denoising and other applications are left for future work.

[5]This version [69] was observed to typically provide similar quality of results in our experiments as [68].

Fig. 2. The 512×512 standard images used in our experiments. The images are Barbara, Boat, Couple, Hill, and Lena.

and Hill, shown in Fig. 2. The SOUP-DIL Algorithm for (P1) was then used to learn a 64×256 overcomplete dictionary, with $\lambda = 69$. The algorithm is initialized as mentioned in Section III-A. Specifically, the initial estimate for C is an all-zero matrix, and the initial estimate for D is the overcomplete DCT [5], [68].

Fig. 3 illustrates the convergence behavior of SOUP-DIL. The objective in our method (Fig. 3(a)) converged monotonically and quickly over the iterations. Fig. 3(b) shows the normalized sparse representation error and sparsity factor (for C), both expressed as percentages. Both these components of the objective converged quickly for the algorithm, and the NSRE improved by 1 dB beyond the first iteration, indicating the success of the SOUP-DIL approach in representing data using a small number of non-zero coefficients (sparsity factor of 3.14% at convergence).

Importantly, both the quantities $\left\|D^t - D^{t-1}\right\|_F$ (Fig. 3(c)) and $\left\|C^t - C^{t-1}\right\|_F$ (Fig. 3(d)) converge towards 0, as predicted by Theorem 3. This implies that $\left\|Z^t - Z^{t-1}\right\|_F$ with $Z^t = D^t \left(C^t\right)^T$ converges towards zero too. The above results are indicative (are necessary but not sufficient conditions) of the convergence of the entire sequences $\{D^t\}$, $\{C^t\}$, and $\{Z^t\}$ for our algorithm in practice. In contrast, Bao et al. [52] showed that the distance between successive iterates may not converge to zero for popular algorithms such as K-SVD.

C. Sparse Representation of Data

The second experiment worked with the same data as in Section V-B and learned dictionaries of size 64×256 for various choices of the parameter λ in (P1) (i.e., corresponding to a variety of solution sparsity levels). We measure the quality of the learned data approximations DC^T using the NSRE metric. We compare the NSRE values achieved by our algorithm to those achieved by the K-SVD dictionary learning scheme [28] for the same data. The K-SVD method was executed for several choices of sparsity (number of non-zeros) of the columns of C^T. The λ values for (P1) were chosen so as to achieve similar average column sparsity levels

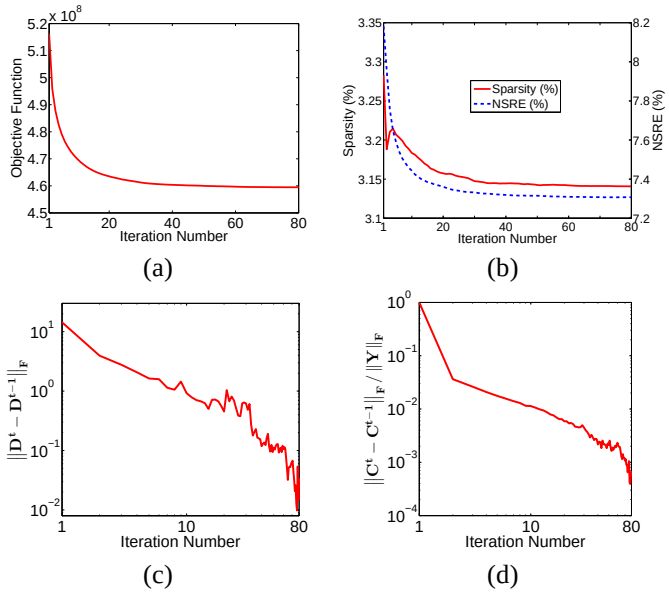

Fig. 3. Convergence behavior of the SOUP-DIL Algorithm: (a) Objective function; (b) Normalized sparse representation error (percentage) and sparsity factor of C ($\sum_{j=1}^{J} \left(\|c_j\|_0 / nN\right)$ – expressed as a percentage); (c) changes between successive D iterates ($\left\|D^t - D^{t-1}\right\|_F$); and (d) normalized changes between successive C iterates ($\left\|C^t - C^{t-1}\right\|_F / \|Y\|_F$).

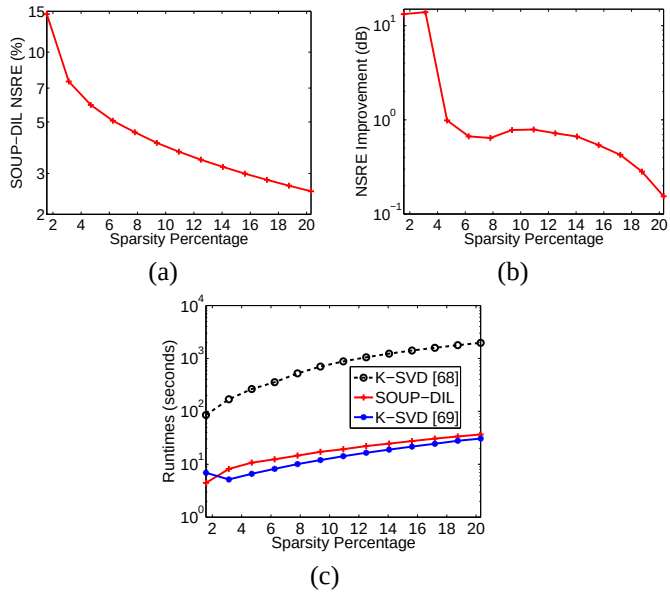

Fig. 4. Behavior of the SOUP-DIL Algorithm and K-SVD [28] at various sparsities (or sparsity factors): (a) NSRE (percentage) for SOUP-DIL; (b) Improvements in NSRE (in decibels) provided by SOUP-DIL over K-SVD; and (c) runtimes for SOUP-DIL, the matlab implementation of K-SVD [28], [68], and the efficient (partial) MEX/C implementation of K-SVD [62], [69].

in C^T as K-SVD. Both the SOUP-DIL Algorithm and K-SVD were initialized with the same overcomplete DCT dictionary in our experiment and both methods ran for 10 iterations.

Fig. 4 shows the behavior of the SOUP-DIL and K-SVD algorithms for average column sparsity levels of C^T ranging from about 2% to 20%. As expected, the NSRE values for the SOUP-DIL Algorithm decreased monotonically (Fig. 4(a)) when the average column sparsity level in C^T increased

(i.e., as λ decreased). Importantly, the SOUP-DIL Algorithm provided better data representations (Fig. 4(b)) than K-SVD at the various tested sparsity levels. Large improvements of about 14 dB and 1 dB are observed at low and mid sparsity levels, respectively.

Finally, Fig. 4(c) compares the runtimes of our method to those of the unoptimized Matlab implementation of K-SVD [68] as well as the efficient (partial) MEX/C implementation [69] of K-SVD demonstrating large speedups of 40-50 times for SOUP-DIL over the first K-SVD implementation (at most sparsities), while the runtimes for the SOUP-DIL method are about the same as those of the second K-SVD implementation. Since these results were obtained using only an unoptimized Matlab implementation of SOUP-DIL, we expect significant speed-ups for our scheme with code optimization or C/C++ implementations.

D. Image Denoising

In image denoising, the goal is to recover an estimate of an image $x \in \mathbb{R}^M$ (2D image represented as a vector) from its corrupted measurements $y = x + h$, where h is the noise. For example, the entries of h may be i.i.d. Gaussian with zero mean and standard deviation σ.

To perform image denoising using (P1), we first extract all the overlapping patches (with maximum overlap) of the noisy image y, and construct the training set $Y \in \mathbb{R}^{n \times N}$ as a matrix whose columns are those noisy patches. We then use Problem (P1) to learn a dictionary and sparse codes for Y. To obtain the denoised image estimate, we then solve the following least squares problem, where $\hat{D}$ and $\hat{\alpha}_j$ denote the learned dictionary and patch sparse codes obtained from the noisy patches, and P_j is an operator that extracts a patch as a vector:

$$\min_{x} \sum_{j=1}^{N} \left\| P_j x - \hat{D}\hat{\alpha}_j \right\|_2^2 + \nu \left\| x - y \right\|_2^2 \qquad (22)$$

The optimal $\hat{x}$ is obtained [5] by summing together the denoised patch estimates $\hat{D}\hat{\alpha}_j$ at their respective 2D locations, and computing a weighted average between this result and the noisy image.

K-SVD based denoising [5] involves a similar methodology as described above for (P1), but differs in the dictionary learning procedure, where the ℓ_0 "norms" of the sparse codes are minimized so that a fitting constraint or error constraint of $\left\| P_j y - \hat{D}\hat{\alpha}_j \right\|_2^2 \leq \epsilon$ is met for representing the noisy patches. In particular, when the noise is i.i.d. Gaussian, $\epsilon = nC^2\sigma^2$ is used, with $C > 1$ a constant. Such a constraint serves as a strong prior [5], [6], and is a key reason for the denoising capability of K-SVD. Hence, in our method based on (P1), we set $\lambda \propto \sigma$ (noise standard deviation) in (P1) during learning, and once the dictionary $\hat{D}$ is learned from noisy patches, we estimate the patch sparse codes $\hat{\alpha}_j$ using a single pass (over the noisy patches) of orthogonal matching pursuit (OMP) by employing an error constraint criterion like in K-SVD. Although we only use information on the noise statistics in a sub-optimal way for (P1), we still show good denoising performance (vis-a-vis K-SVD) in the following with this approach. Note that for both our method and for K-SVD [5], [68], the means of the patches are removed prior to learning and added back afterwards to obtain the denoised estimates.

Image	σ	Noisy	O-DCT	K-SVD	SOUP-DIL
Couple	5	34.16	37.25	37.29	37.28
	10	28.11	33.40	33.49	33.50
	20	22.11	29.71	30.01	29.99
	25	20.17	28.53	28.88	28.92
	30	18.58	27.53	27.88	27.97
	100	8.13	22.59	22.58	22.71
Barbara	5	34.15	37.94	38.08	38.04
	10	28.14	33.96	34.43	34.37
	20	22.13	29.95	30.83	30.79
	25	20.17	28.68	29.63	29.64
	30	18.59	27.62	28.54	28.63
	100	8.11	21.87	21.87	21.97
Boat	5	34.15	37.09	37.21	37.16
	10	28.13	33.43	33.62	33.60
	20	22.10	29.92	30.36	30.37
	25	20.17	28.79	29.28	29.30
	30	18.60	27.93	28.41	28.43
	100	8.13	22.79	22.81	22.96
Hill	5	34.15	37.02	37.08	37.05
	10	28.14	33.26	33.45	33.44
	20	22.10	29.85	30.17	30.20
	25	20.18	28.89	29.23	29.31
	30	18.57	28.14	28.43	28.56
	100	8.16	24.00	23.98	24.03
Lena	5	34.16	38.52	38.62	38.55
	10	28.12	35.30	35.48	35.47
	20	22.11	32.02	32.40	32.40
	25	20.18	30.89	31.32	31.32
	30	18.59	29.98	30.41	30.46
	100	8.14	24.45	24.51	24.63
Avg.	5	34.16	37.56	37.66	37.61
	10	28.13	33.87	34.09	34.07
	20	22.11	30.29	30.75	30.75
	25	20.17	29.16	29.67	29.70
	30	18.58	28.24	28.74	28.81
	100	8.13	23.14	23.15	23.26

TABLE I
PSNR VALUES IN DECIBELS FOR DENOISING WITH OVERCOMPLETE DCT (O-DCT), K-SVD [5], AND SOUP-DIL. ALL DICTIONARIES HAVE SIZE 64×256. THE PSNR VALUES OF THE NOISY IMAGES (NOISY) ARE ALSO SHOWN. THE LAST BLOCK OF ENTRIES IN THE TABLE CORRESPONDS TO THE AVERAGE (OVER THE IMAGES) PSNR VALUES.

For the denoising experiments, we work with the images Couple, Barbara, Boat, Hill, and Lena in Fig. 2, and simulate i.i.d. Gaussian noise at six different noise levels ($\sigma = 5$, 10, 20, 25, 30, 100) for each of the images. We then denoise these images using both the SOUP-DIL denoising method outlined above and K-SVD [5], [68]. SOUP-DIL denoising uses patches of size 8×8, a 64×256 overcomplete dictionary, $\lambda = 5\sigma$, $\nu = 20/\sigma$ (in (22)), an overcomplete DCT initial dictionary in learning, and 10 iterations of the learning algorithm. These parameter settings[6] were found to work well for our method.

Table I lists the denoising PSNRs obtained by the SOUP-

[6]The built-in parameter settings in the K-SVD denoising implementation [5], [68] are very similar except that K-SVD uses a subset of the patches for training. We found that the denoising performance of K-SVD changes by at most few hundredths of a dB when all the patches are used in learning. However, this setting also leads to increased runtimes.

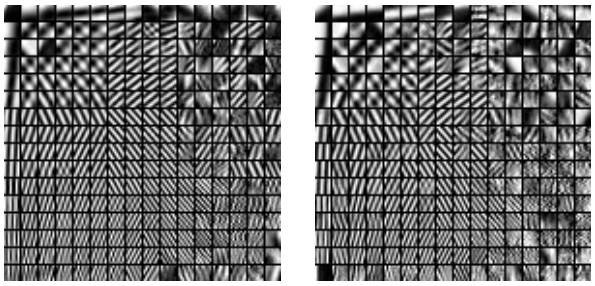

(a) (b)

Fig. 5. Learned dictionaries for Barbara at $\sigma = 20$: (a) SOUP dictionary; and (b) K-SVD dictionary. The columns of the dictionaries are shown as patches.

DIL denoising method, along with the PSNRs obtained by K-SVD. For comparison, we also list the denoising PSNRs obtained by employing the overcomplete DCT dictionary for denoising. In the latter case, the same strategy (and parameter settings [68]) as used by K-SVD based denoising is adopted but while skipping the learning process.

The various methods denoise about the same in Table I for a low noise level of $\sigma = 5$. For $\sigma > 5$, the SOUP-DIL method denoises about 0.4 dB better on the average than the overcomplete DCT. While the SOUP-DIL method and K-SVD denoise about the same at low and mid noise levels, the SOUP-DIL scheme performs about 0.1 dB better on average at higher noise levels such as $\sigma = 30$ or 100 in Table I. Importantly, SOUP-DIL based denoising is highly efficient and the learning procedure has good convergence properties. We have observed similar effects as in Section V-C for the runtimes of SOUP-DIL denoising vis-a-vis the unoptimized Matlab [68] or the efficient (partial MEX/C) [69] implementations of the K-SVD approach.

Fig. 5 shows the dictionaries learned from noisy patches using the SOUP-DIL method and K-SVD for the image Barbara at $\sigma = 20$. Both dictionaries show frequency and textural features that are specific to the image Barbara. By learning such image-specific features, the SOUP-DIL method (and K-SVD) easily outperforms fixed dictionaries such as the overcomplete DCT in denoising.

VI. Conclusions

This paper proposed a fast method for synthesis dictionary learning. The training data set is approximated by a sum of sparse rank-one matrices. The proposed formulation learns the left and right singular vectors of these rank-one matrices using an ℓ_0 penalty to enforce sparsity. We adopted a block coordinate descent method to efficiently update the unknown sparse codes and dictionary atoms in the problem. In particular, the sparse coding step in our algorithm is performed cheaply by truncated hard-thresholding, and the dictionary atom update involves a single sparse vector-matrix product that is computed efficiently. A convergence analysis was presented for the proposed block coordinate descent algorithm for a highly non-convex problem. The proposed approach had comparable or superior performance and significant speed-ups over the classical K-SVD method [5], [28] in sparse signal representation and image denoising. The usefulness of our method in other applications such as blind compressed sensing [70] or other inverse problems in imaging merits further study. Extensions of the proposed method for online learning [8] are also of potential interest.

Appendix A
Proof of Theorem 1

First, we discuss the convergence of the objective sequence. At every iteration t and inner iteration j in Fig. 1, we solve the sparse coding and dictionary atom update subproblems exactly. Thus, we have the following two inequalities:

$$f\left(c_1^t, ..., c_{j-1}^t, c_j^t, c_{j+1}^{t-1}, ..., c_J^{t-1}, d_1^t, ..., d_{j-1}^t, d_j^{t-1}, ..., d_J^{t-1}\right) \leq f\left(c_1^t, ..., c_{j-1}^t, c_j^{t-1}, c_{j+1}^{t-1}, ..., c_J^{t-1}, d_1^t, ..., d_{j-1}^t, d_j^{t-1}, ..., d_J^{t-1}\right)$$
$$f\left(c_1^t, ..., c_j^t, c_{j+1}^{t-1}, ..., c_J^{t-1}, d_1^t, ..., d_{j-1}^t, d_j^t, d_{j+1}^{t-1}, ..., d_J^{t-1}\right) \leq f\left(c_1^t, ..., c_j^t, c_{j+1}^{t-1}, ..., c_J^{t-1}, d_1^t, ..., d_{j-1}^t, d_j^{t-1}, d_{j+1}^{t-1}, ..., d_J^{t-1}\right)$$

Combining these inequalities for the tth iteration and for each $1 \leq j \leq J$ yields $f(C^t, D^t) \leq f(C^{t-1}, D^{t-1})$. Thus, the sequence $\{f(C^t, D^t)\}$ is monotone decreasing. Because it is lower bounded (by 0), it converges to a finite value $f^* = f^*(C^0, D^0)$ (that may depend on the initial conditions).

Next, the boundedness of the $\{D^t\}$ and $\{C^t\}$ sequences is obvious from the constraints in Problem (P1), so the accumulation points of the iterates form a non-empty and compact set.

Finally, we show that each accumulation point achieves the same value f^* of the objective. Consider a subsequence $\{C^{q_t}, D^{q_t}\}$ of the iterate sequence that converges to the accumulation point (C^*, D^*). Because $\left\|d_j^{q_t}\right\|_2 = 1$ and $\left\|c_j^{q_t}\right\|_\infty \leq L$ for $1 \leq j \leq J$ and every t, therefore, due to the continuity of the norms, we also have $\left\|d_j^*\right\|_2 = 1$ and $\left\|c_j^*\right\|_\infty \leq L$ for all j. Thus, the barrier functions in (21) are inactive for the subsequence and its limit, i.e.,

$$\chi(d_j^*) = 0,\ \psi(c_j^*) = 0\ \forall\, j. \tag{23}$$

The formula for c_j^t in (18) implies that it cannot have any non-zero entry of magnitude less than λ. Because $\left\{c_j^{q_t}\right\}$ converges to c_j^* entry-wise, therefore, for any entry of c_j^* that is zero, the corresponding entry in $c_j^{q_t} = 0$ for large enough t. Clearly, wherever c_j^* is non-zero, the corresponding entry in $c_j^{q_t}$ stays non-zero for large enough t values. Thus,

$$\lim_{t\to\infty} \left\|c_j^{q_t}\right\|_0 = \left\|c_j^*\right\|_0\ \forall\, j \tag{24}$$

and the convergence in (24) happens in a finite number of iterations. We then easily have the following result:

$$\lim_{t\to\infty} f(C^{q_t}, D^{q_t}) = \lim_{t\to\infty} \left\|Y - D^{q_t}\left(C^{q_t}\right)^T\right\|_F^2 + \lambda^2 \sum_{j=1}^{J} \lim_{t\to\infty} \left\|c_j^{q_t}\right\|_0 = \lambda^2 \sum_{j=1}^{J} \left\|c_j^*\right\|_0 + \left\|Y - D^*\left(C^*\right)^T\right\|_F^2$$

The right hand side above coincides with $f(C^*, D^*)$. Since the objective sequence converges to f^*, therefore, $f(C^*, D^*) = \lim_{t\to\infty} f(C^{q_t}, D^{q_t}) = f^*$. ■

APPENDIX B
PROOF OF THEOREM 2

First, consider a subsequence $\{c^{q_t}, d^{q_t}\}$ of the iterate sequence that converges to the accumulation point (c^*, d^*). Because of the optimality of d^{q_t} in the dictionary atom update step (11) of the q_tth iteration, we have the following inequality for all $d \in \mathbb{R}^n$ with $\|d\|_2 = 1$:

$$\left\|Y - d^{q_t}\left(c^{q_t}\right)^T\right\|_F^2 + \lambda^2 \left\|c^{q_t}\right\|_0 \leq \left\|Y - d\left(c^{q_t}\right)^T\right\|_F^2 + \lambda^2 \left\|c^{q_t}\right\|_0. \quad (25)$$

Taking the limit $t \to \infty$ in (25) and using (24) leads to the following result that holds for each feasible d:

$$\left\|Y - d^*\left(c^*\right)^T\right\|_F^2 + \lambda^2 \left\|c^*\right\|_0 \leq \left\|Y - d\left(c^*\right)^T\right\|_F^2 + \lambda^2 \left\|c^*\right\|_0.$$

The result implies that d^* (which has unit ℓ_2 norm) is in fact a global minimizer of $f(c^*, d)$ with respect to d, i.e., $0 \in \partial f_d(c^*, d^*)$.

Next, consider a convergent subsequence $\left\{c^{q_{n_t}+1}\right\}$ of the bounded sequence $\left\{c^{q_t+1}\right\}$ that has a limit c^{**}. Because of the optimality of $c^{q_{n_t}+1}$ in the sparse coding step (2) of iteration $q_{n_t}+1$, we have the following inequality for all $c \in \mathbb{R}^N$ and for $f(c,d) = \left\|Y - dc^T\right\|_F^2 + \lambda^2 \|c\|_0 + \chi(d) + \psi(c)$:

$$f(c^{q_{n_t}+1}, d^{q_{n_t}}) \leq f(c, d^{q_{n_t}}) \quad (26)$$

As in (25), taking the limit $t \to \infty$ above and using (24) and (23) yields the following result that holds for each c:

$$f(c^{**}, d^*) \leq f(c, d^*) \quad (27)$$

This implies that c^{**} (that has $\|c^{**}\|_\infty \leq L$) is a global minimizer of $f(c, d^*)$ with respect to c achieving the objective value $f(c^{**}, d^*)$.

Since $f(c^{q_{n_t}+1}, d^{q_{n_t}+1}) \leq f(c^{q_{n_t}+1}, d^{q_{n_t}}) \leq f(c^{q_{n_t}}, c^{q_{n_t}})$, and by Theorem 1, the first and last terms in this inequality converge to f^*, we therefore have that $f(c^{**}, d^*) = f^*$ too. By Theorem 1, $f(c^*, d^*) = f^*$ holds for the accumulation point (c^*, d^*) of the subsequence $\{c^{q_t}, d^{q_t}\}$, implying $f(c^*, d^*) = f(c^{**}, d^*)$. Combining this with (27), it is clear that c^* is also a global minimizer of $f(c, d^*)$ with respect to c, i.e., $0 \in \partial f_c(c^*, d^*)$.

Finally, (see Proposition 3 in [71]) the subdifferential ∂f at (c^*, d^*) satisfies $\partial f(c^*, d^*) = \partial f_c(c^*, d^*) \times \partial f_d(c^*, d^*)$. Since $0 \in \partial f_d(c^*, d^*)$ and $0 \in \partial f_c(c^*, d^*)$, we therefore have that $0 \in \partial f(c^*, d^*)$. Thus, when $J = 1$ in the SOUP-DIL Algorithm, each accumulation point in the algorithm is a critical point of the objective f. ■

APPENDIX C
PROOF OF THEOREM 3

A. Critical Point Property

Consider a convergent subsequence $\{C^{q_t}, D^{q_t}\}$ of the iterate sequence in the SOUP-DIL Algorithm that converges to (C^*, D^*). Let $\left\{C^{q_{n_t}+1}, D^{q_{n_t}+1}\right\}$ be a convergent subsequence of the bounded $\left\{C^{q_t+1}, D^{q_t+1}\right\}$, with limit (C^{**}, D^{**}). For each iteration t and inner iteration j in the algorithm, define the matrix $E_j^t \triangleq Y - \sum_{k<j} d_k^t \left(c_k^t\right)^T - \sum_{k>j} d_k^{t-1}\left(c_k^{t-1}\right)^T$. For the accumulation point (C^*, D^*), define $E_j^* \triangleq Y - D^*\left(C^*\right)^T + d_j^*\left(c_j^*\right)^T$. In this proof, for simplicity, we denote the objective f (21) in the jth sparse coding step of iteration t (Fig. 1) as

$$f\left(E_j^t, c_j, d_j^{t-1}\right) \triangleq \left\|E_j^t - d_j^{t-1}c_j^T\right\|_F^2 + \lambda^2 \sum_{k<j} \left\|c_k^t\right\|_0 + \lambda^2 \sum_{k>j} \left\|c_k^{t-1}\right\|_0 + \lambda^2 \|c_j\|_0 + \psi(c_j) \quad (28)$$

All but the jth atom and sparse vector c_j are represented via E_j^t on the left hand side in this notation. The objective that is minimized in the dictionary atom update step is similarly denoted as $f\left(E_j^t, c_j^t, d_j\right)$ with

$$f\left(E_j^t, c_j^t, d_j\right) \triangleq \left\|E_j^t - d_j\left(c_j^t\right)^T\right\|_F^2 + \lambda^2 \sum_{k\leq j} \left\|c_k^t\right\|_0 + \lambda^2 \sum_{k>j} \left\|c_k^{t-1}\right\|_0 + \chi(d_j). \quad (29)$$

Finally, the functions $f\left(E_j^*, c_j, d_j^*\right)$ and $f\left(E_j^*, c_j^*, d_j\right)$ are defined in a similar way with respect to the accumulation point (C^*, D^*).

To establish the critical point property of (C^*, D^*), we first show the partial global optimality of each column of the matrices C^* and D^* for f. By partial global optimality, we mean that each column of C^* or D^* is a global minimizer of f, when all other variables are kept fixed to the values in (C^*, D^*). First, for $j = 1$ and iteration $q_{n_t}+1$, we have the following result for the sparse coding step for all $c_1 \in \mathbb{R}^N$:

$$f\left(E_1^{q_{n_t}+1}, c_1^{q_{n_t}+1}, d_1^{q_{n_t}}\right) \leq f\left(E_1^{q_{n_t}+1}, c_1, d_1^{q_{n_t}}\right) \quad (30)$$

Taking the limit $t \to \infty$ above and using (24) to obtain limits of ℓ_0 terms in the cost (28), and using (23), and the fact that $E_1^{q_{n_t}+1} \to E_1^*$, we have

$$f\left(E_1^*, c_1^{**}, d_1^*\right) \leq f\left(E_1^*, c_1, d_1^*\right) \forall\, c_1 \in \mathbb{R}^N. \quad (31)$$

This means that c_1^{**} is a minimizer of f with all other variables fixed to their values in (C^*, D^*). Because of the (uniqueness) assumption in the theorem, we have

$$c_1^{**} = \arg\min_{c_1} f\left(E_1^*, c_1, d_1^*\right) \quad (32)$$

Furthermore, because of the equivalence of accumulation points, $f\left(E_1^*, c_1^{**}, d_1^*\right) = f\left(E_1^*, c_1^*, d_1^*\right) = f^*$ holds. This result together with (32) implies that $c_1^{**} = c_1^*$ and thus

$$c_1^* = \arg\min_{c_1} f\left(E_1^*, c_1, d_1^*\right) \quad (33)$$

Therefore, c_1^* is a partial global minimizer of f.

Next, for the first dictionary atom update step ($j = 1$) in iteration $q_{n_t}+1$, we have the following for all $d_1 \in \mathbb{R}^n$:

$$f\left(E_1^{q_{n_t}+1}, c_1^{q_{n_t}+1}, d_1^{q_{n_t}+1}\right) \leq f\left(E_1^{q_{n_t}+1}, c_1^{q_{n_t}+1}, d_1\right) \quad (34)$$

Just like in (30), upon taking the limit $t \to \infty$ above and using $c_1^{**} = c_1^*$, we get

$$f\left(E_1^*, c_1^*, d_1^{**}\right) \leq f\left(E_1^*, c_1^*, d_1\right) \forall\, d_1 \in \mathbb{R}^n. \quad (35)$$

Thus, d_1^{**} is a minimizer of the cost $f\left(E_1^*, c_1^*, d_1\right)$ with respect to d_1. Because of the equivalence of accumulation points, we

have $f\left(E_1^*, c_1^*, d_1^{**}\right) = f\left(E_1^*, c_1^*, d_1^*\right) = f^*$. This implies that d_1^* is also a partial global minimizer of f in (35) satisfying

$$d_1^* \in \arg\min_{d_1} f\left(E_1^*, c_1^*, d_1\right) \tag{36}$$

By Proposition 2, the minimizer of the dictionary atom update cost is unique as long as the corresponding sparse code (in (33)) is non-zero. Therefore, $d_1^{**} = d_1^*$ is the *unique* minimizer in (36), except when $c_1^* = 0$.

When $c_1^* = 0$, we use (24) to conclude that $c_1^{q_t} = 0$ for all sufficiently large t values. Since $c_1^{**} = c_1^*$, we must also have that $c_1^{q_{n_t}+1} = 0$ for all large enough t. Therefore, for sufficiently large t, $d_1^{q_t}$ and $d_1^{q_{n_t}+1}$ are the minimizers of dictionary atom update steps (11), wherein the corresponding sparse coefficients $c_1^{q_t}$ and $c_1^{q_{n_t}+1}$ are zero, implying that $d_1^{q_t} = d_1^{q_{n_t}+1} = v_1$ (Fig. 1) for all sufficiently large t. Thus, the limits satisfy $d_1^* = d_1^{**} = v_1$. Therefore, $d_1^{**} = d_1^*$ holds, even when $c_1^* = 0$. Therefore, for $j = 1$,

$$d_1^{**} = d_1^*, \;\; c_1^{**} = c_1^*. \tag{37}$$

Next, we repeat the above procedure by considering first the sparse coding step and then the dictionary update step for $j = 2$ and iteration $q_{n_t} + 1$. For $j = 2$, we consider the matrix

$$E_2^{q_{n_t}+1} = Y - \sum_{k>2} d_k^{q_{n_t}} \left(c_k^{q_{n_t}}\right)^T - d_1^{q_{n_t}+1} \left(c_1^{q_{n_t}+1}\right)^T$$

It follows from (37) that $E_2^{q_{n_t}+1} \to E_2^*$ as $t \to \infty$. Then, by repeating the steps (30) - (37), we can easily show that c_2^* and d_2^* are each partial global minimizers of f when all other variables are fixed to their values in (C^*, D^*). Moreover, $c_2^{**} = c_2^*$ and $d_2^{**} = d_2^*$. Similar such arguments can be made sequentially for all other values of j until $j = J$.

Finally, the partial global optimality of each column of C^* and D^* for the objective f implies that $0 \in \partial f\left(C^*, D^*\right)$, i.e., (C^*, D^*) is a critical point of the function f. ■

B. Convergence of the Difference between Successive Iterates

Consider the sequence $\{a^t\}$ whose elements are $a^t \triangleq \left\|D^t - D^{t-1}\right\|_F$. Clearly, this sequence is bounded. We will show that every convergent subsequence of this sequence converges to zero, thereby implying that zero is the only accumulation point, i.e., $\{a^t\}$ converges to 0. A similar argument establishes that $\left\|C^t - C^{t-1}\right\|_F \to 0$ as $t \to \infty$.

Consider a convergent subsequence $\{a^{q_t}\}$ of the sequence $\{a^t\}$. The bounded sequence $\left\{\left(C^{q_t-1}, D^{q_t-1}, C^{q_t}, D^{q_t}\right)\right\}$ (whose elements are formed by pairing successive elements of the iterate sequence) must have a convergent subsequence $\left\{\left(C^{q_{n_t}-1}, D^{q_{n_t}-1}, C^{q_{n_t}}, D^{q_{n_t}}\right)\right\}$ that converges to a point $(C^*, D^*, C^{**}, D^{**})$. Based on the results in Appendix C-A, we then have $d_j^{**} = d_j^*$ (and $c_j^{**} = c_j^*$) for each $1 \le j \le J$, or

$$D^{**} = D^*. \tag{38}$$

Thus, clearly $a^{q_{n_t}} \to 0$ as $t \to \infty$. Since, $\{a^{q_{n_t}}\}$ is a subsequence of the convergent $\{a^{q_t}\}$, we must have that $\{a^{q_t}\}$ converges to zero too. Thus, we have shown that zero is the limit of any arbitrary convergent subsequence of $\{a^t\}$. ■

REFERENCES

[1] M. W. Marcellin, M. J. Gormish, A. Bilgin, and M. P. Boliek, "An overview of JPEG-2000," in *Proc. Data Compression Conf.*, 2000, pp. 523–541.

[2] M. Elad, P. Milanfar, and R. Rubinstein, "Analysis versus synthesis in signal priors," *Inverse Problems*, vol. 23, no. 3, pp. 947–968, 2007.

[3] W. K. Pratt, J. Kane, and H. C. Andrews, "Hadamard transform image coding," *Proc. IEEE*, vol. 57, no. 1, pp. 58–68, 1969.

[4] S. Ravishankar and Y. Bresler, "Learning sparsifying transforms," *IEEE Trans. Signal Process.*, vol. 61, no. 5, pp. 1072–1086, 2013.

[5] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, 2006.

[6] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. on Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.

[7] M. Protter and M. Elad, "Image sequence denoising via sparse and redundant representations," *IEEE Trans. on Image Processing*, vol. 18, no. 1, pp. 27–36, 2009.

[8] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, 2010.

[9] S. Ravishankar and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, no. 5, pp. 1028–1041, 2011.

[10] R. Rubinstein, T. Peleg, and M. Elad, "Analysis K-SVD: A dictionary-learning algorithm for the analysis sparse model," *IEEE Transactions on Signal Processing*, vol. 61, no. 3, pp. 661–677, 2013.

[11] S. Ravishankar and Y. Bresler, "Learning doubly sparse transforms for images," *IEEE Trans. Image Process.*, vol. 22, no. 12, pp. 4598–4612, 2013.

[12] B. Wen, S. Ravishankar, and Y. Bresler, "Structured overcomplete sparsifying transform learning with convergence guarantees and applications," *International Journal of Computer Vision*, vol. 114, no. 2-3, pp. 137–167, 2015.

[13] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, 2009.

[14] R. Vidal, "Subspace clustering," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 52–68, 2011.

[15] E. Elhamifar and R. Vidal, "Sparsity in unions of subspaces for classification and clustering of high-dimensional data," in *49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2011, pp. 1085–1089.

[16] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.

[17] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *Journal of Constructive Approximation*, vol. 13, no. 1, pp. 57–98, 1997.

[18] Y. Pati, R. Rezaiifar, and P. Krishnaprasad, "Orthogonal matching pursuit : recursive function approximation with applications to wavelet decomposition," in *Asilomar Conf. on Signals, Systems and Comput.*, 1993, pp. 40–44 vol.1.

[19] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.

[20] I. F. Gorodnitsky, J. George, and B. D. Rao, "Neuromagnetic source imaging with FOCUSS: A recursive weighted minimum norm algorithm," *Electrocephalography and Clinical Neurophysiology*, vol. 95, pp. 231–251, 1995.

[21] G. Harikumar and Y. Bresler, "A new algorithm for computing sparse solutions to linear inverse problems," in *ICASSP*, may 1996, pp. 1331–1334.

[22] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.

[23] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, "Least angle regression," *Annals of Statistics*, vol. 32, pp. 407–499, 2004.

[24] D. Needell and J. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301 – 321, 2009.

[25] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Information Theory*, vol. 55, no. 5, pp. 2230–2249, 2009.

[26] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.

[27] K. Engan, S. Aase, and J. Hakon-Husoy, "Method of optimal directions for frame design," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1999, pp. 2443–2446.

[28] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on signal processing*, vol. 54, no. 11, pp. 4311–4322, 2006.

[29] M. Yaghoobi, T. Blumensath, and M. Davies, "Dictionary learning for sparse approximations with the majorization method," *IEEE Transaction on Signal Processing*, vol. 57, no. 6, pp. 2178–2191, 2009.

[30] O. Bryt and M. Elad, "Compression of facial images using the k-svd algorithm," *Journal of Visual Communication and Image Representation*, vol. 19, no. 4, pp. 270–282, 2008.

[31] M. Aharon and M. Elad, "Sparse and redundant modeling of image content using an image-signature-dictionary," *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 228–247, 2008.

[32] J. Mairal, G. Sapiro, and M. Elad, "Learning multiscale sparse representations for image and video restoration," *SIAM Multiscale Modeling and Simulation*, vol. 7, no. 1, pp. 214–241, 2008.

[33] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *IEEE International Conference on Computer Vision*, Sept 2009, pp. 2272–2279.

[34] X. Lu, Y. Yuan, and P. Yan, "Alternatively constrained dictionary learning for image superresolution," *IEEE Transactions on Cybernetics*, vol. 44, no. 3, pp. 366–377, 2014.

[35] I. Ramirez, P. Sprechmann, and G. Sapiro, "Classification and clustering via dictionary learning with structured incoherence and shared features," in *Proc. IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2010*, 2010, pp. 3501–3508.

[36] S. Kong and D. Wang, "A dictionary learning approach for classification: Separating the particularity and the commonality," in *Proceedings of the 12th European Conference on Computer Vision*, 2012, pp. 186–199.

[37] Z. Jiang, Z. Lin, and L. S. Davis, "Label consistent k-svd: Learning a discriminative dictionary for recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 11, pp. 2651–2664, 2013.

[38] H. Y. Liao and G. Sapiro, "Sparse representations for limited data tomography," in *Proc. IEEE International Symposium on Biomedical Imaging (ISBI)*, 2008, pp. 1375–1378.

[39] S. Ravishankar and Y. Bresler, "Multiscale dictionary learning for MRI," in *Proc. ISMRM*, 2011, p. 2830.

[40] Y. Wang, Y. Zhou, and L. Ying, "Undersampled dynamic magnetic resonance imaging using patch-based spatiotemporal dictionaries," in *2013 IEEE 10th International Symposium on Biomedical Imaging (ISBI)*, April 2013, pp. 294–297.

[41] Y. Huang, J. Paisley, Q. Lin, X. Ding, X. Fu, and X. P. Zhang, "Bayesian nonparametric dictionary learning for compressed sensing mri," *IEEE Trans. Image Process.*, vol. 23, no. 12, pp. 5007–5019, 2014.

[42] R. Gribonval and K. Schnass, "Dictionary identification–sparse matrix-factorization via l_1 -minimization," *IEEE Trans. Inform. Theory*, vol. 56, no. 7, pp. 3523–3539, 2010.

[43] D. Barchiesi and M. D. Plumbley, "Learning incoherent dictionaries for sparse approximation using iterative projections and rotations," *IEEE Transactions on Signal Processing*, vol. 61, no. 8, pp. 2055–2065, 2013.

[44] R. Rubinstein, M. Zibulevsky, and M. Elad, "Double sparsity: Learning sparse dictionaries for sparse signal approximation," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1553–1564, 2010.

[45] K. Skretting and K. Engan, "Recursive least squares dictionary learning algorithm," *IEEE Transactions on Signal Processing*, vol. 58, no. 4, pp. 2121–2130, 2010.

[46] B. Ophir, M. Lustig, and M. Elad, "Multi-scale dictionary learning using wavelets," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 1014–1024, 2011.

[47] S. K. Sahoo and A. Makur, "Dictionary training for sparse representation as generalization of k-means clustering," *IEEE Signal Processing Letters*, vol. 20, no. 6, pp. 587–590, June 2013.

[48] L. N. Smith and M. Elad, "Improving dictionary learning: Multiple dictionary updates and coefficient reuse," *IEEE Signal Processing Letters*, vol. 20, no. 1, pp. 79–82, Jan 2013.

[49] M. Sadeghi, M. Babaie-Zadeh, and C. Jutten, "Dictionary learning for sparse representation: A novel approach," *IEEE Signal Processing Letters*, vol. 20, no. 12, pp. 1195–1198, Dec 2013.

[50] A.-K. Seghouane and M. Hanif, "A sequential dictionary learning algorithm with enforced sparsity," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 3876–3880.

[51] Y. Naderahmadian, S. Beheshti, and M. Tinati, "Correlation based online dictionary learning algorithm," *IEEE Transactions on Signal Processing*, 2015.

[52] C. Bao, H. Ji, Y. Quan, and Z. Shen, "L0 norm based dictionary learning by proximal methods with global convergence," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 3858–3865.

[53] A. Rakotomamonjy, "Direct optimization of the dictionary learning problem," *IEEE Transactions on Signal Processing*, vol. 61, no. 22, pp. 5495–5506, 2013.

[54] S. Hawe, M. Seibert, and M. Kleinsteuber, "Separable dictionary learning," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 438–445.

[55] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proceedings of the 25th Annual Conference on Learning Theory*, 2012, pp. 37.1–37.18.

[56] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning sparsely used overcomplete dictionaries via alternating minimization," 2013, preprint: http://arxiv.org/abs/1310.7991.

[57] S. Arora, R. Ge, and A. Moitra, "New algorithms for learning incoherent and overcomplete dictionaries," 2013, preprint: http://arxiv.org/pdf/1308.6273v5.pdf.

[58] Y. Xu and W. Yin, "A fast patch-dictionary method for whole-image recovery," 2013, UCLA CAM report 13-38. [Online]. Available: ftp://ftp.math.ucla.edu/pub/camreport/cam13-38.pdf

[59] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning sparsely used overcomplete dictionaries," *Journal of Machine Learning Research*, vol. 35, pp. 1–15, 2014.

[60] M. Sadeghi, M. Babaie-Zadeh, and C. Jutten, "Learning overcomplete dictionaries based on atom-by-atom updating," *IEEE Transactions on Signal Processing*, vol. 62, no. 4, pp. 883–891, 2014.

[61] C. Bao, H. Ji, Y. Quan, and Z. Shen, "Dictionary learning for sparse coding: Algorithms and analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015.

[62] R. Rubinstein, M. Zibulevsky, and M. Elad, "Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit," http://www.cs.technion.ac.il/~ronrubin/Publications/KSVD-OMP-v2.pdf, 2008, technion - Computer Science Department - Technical Report.

[63] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *J. Optim. Theory Appl.*, vol. 109, no. 3, pp. 475–494, 2001.

[64] Y. Xu and W. Yin, "A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion," *SIAM Journal on Imaging Sciences*, vol. 6, no. 3, pp. 1758–1789, 2013.

[65] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis, 1st edn.* Springer-Verlag, 1997.

[66] B. S. Mordukhovich, *Variational Analysis and Generalized Differentiation. Vol. I: Basic theory*. Springer-Verlag, 2006.

[67] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3D transform-domain collaborative filtering," *IEEE Trans. on Image Processing*, vol. 16, no. 8, pp. 2080–2095, 2007.

[68] M. Elad, "Michael Elad personal page," http://www.cs.technion.ac.il/~elad/Various/KSVD_Matlab_ToolBox.zip, 2009, [Online; accessed Nov. 2015].

[69] R. Rubinstein, http://www.cs.technion.ac.il/~ronrubin/software.html, 2010, [Online; accessed Nov. 2015].

[70] S. Ravishankar and Y. Bresler, "Efficient blind compressed sensing using sparsifying transforms with convergence guarantees and application to magnetic resonance imaging," *SIAM Journal on Imaging Sciences*, vol. 8, no. 4, pp. 2519–2557, 2015.

[71] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran, "Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-łojasiewicz inequality," *Math. Oper. Res.*, vol. 35, no. 2, pp. 438–457, May 2010.